Synchronization of Spiking Neurons in a Computer Model of the Mammalian Retina

G. T. Kenyon (P-21), J. Theiler (NIS-2), K. R. Moore (NIS-1), B. J. Travis (EES-5), J. Jeffs (P-21), and D. W. Marshak (University of Texas Health Science Center, Houston)

Introduction

In typical man-made imaging devices (*i.e.*, CCDs, video cameras), the output from each pixel changes continuously in proportion to the light intensity at the corresponding location. In the mammalian retina, however, the output neurons, called ganglion cells, do not vary their responses continuously, but rather encode information as sequences of uniform impulses, or spikes. Analysis of spike trains recorded from individual ganglion cells strongly favors the hypothesis that information in the optic nerve is "rate coded." When measuring how a typical mammalian ganglion cell responds to a spot of light presented in its receptive field (the region of the visual space to which that cell is most sensitive), it is usually observed that a different sequence of impulses is generated on each stimulus trial. In general, there exists no precise pattern of interspike intervals that is generated reproducibly from one stimulus presentation to the next. If, however, one estimates the instantaneous firing rate of the cell, which can be done by sorting the spikes generated across many stimulus trials into time bins relative to the onset of the stimulus, a regular—and experimentally reproducible— pattern typically emerges. Thus, electrophysiologists have generally rejected the hypothesis that the mammalian visual system uses a "temporal code," in which the exact sequence of interspike intervals generated by each neuron conveys significant information. Instead, most vision scientists believe, at least at the level of individual neurons, that information is encoded by the instantaneous firing rate. According to the rate-code hypothesis, it is the average number of spikes that arrive within a given time window that is important, while the exact sequence of interspike intervals is ignored. Recent studies that have combined sophisticated analysis using information theory with recordings from individual neurons activated by natural or other complex stimuli have largely reinforced the basic assumptions underlying the rate-code hypothesis.[1] In these studies, most of the information conveyed by individual cells could be accurately modeled by a modified Poisson process, in which spikes are generated randomly with a probability governed by the instantaneous firing rate.

Beyond the Rate-Code

Within the last decade, the dominance of the rate-code hypothesis has been challenged by new experiments in which pairs of visual neurons, both in the retina and in the visual cortex, are monitored simultaneously. These studies have shown that many pairs of visual neurons exhibit a stimulus-selective synchronization.[2,3,4] The stimuli in such experiments typically consist of narrow bars of light projected against a uniform background. Pairs of neurons, which may be separated by many degrees of visual angle, are found to fire synchronously when stimulated by a single, long bar of light, while the same pair of neurons fires asynchronously when stimulated by two separate light bars. Extrapolating these results to the case of more complex images, one can imagine how the different components of an image could be segmented by separately synchronized groups of neurons.

Thus, in the cartoon in Figure 1, neurons responding to the same object, here represented by a banana, fire synchronously, while neurons responding to different objects, represented by an apple and an orange, fire asynchronously. A number of investigators have suggested that this behavior might serve as a mechanism for dynamically binding together elements of the visual scene that are functionally related, thus providing a possible neurobiological underpinning to the principles of perceptual grouping identified by Gestalt psychologists.

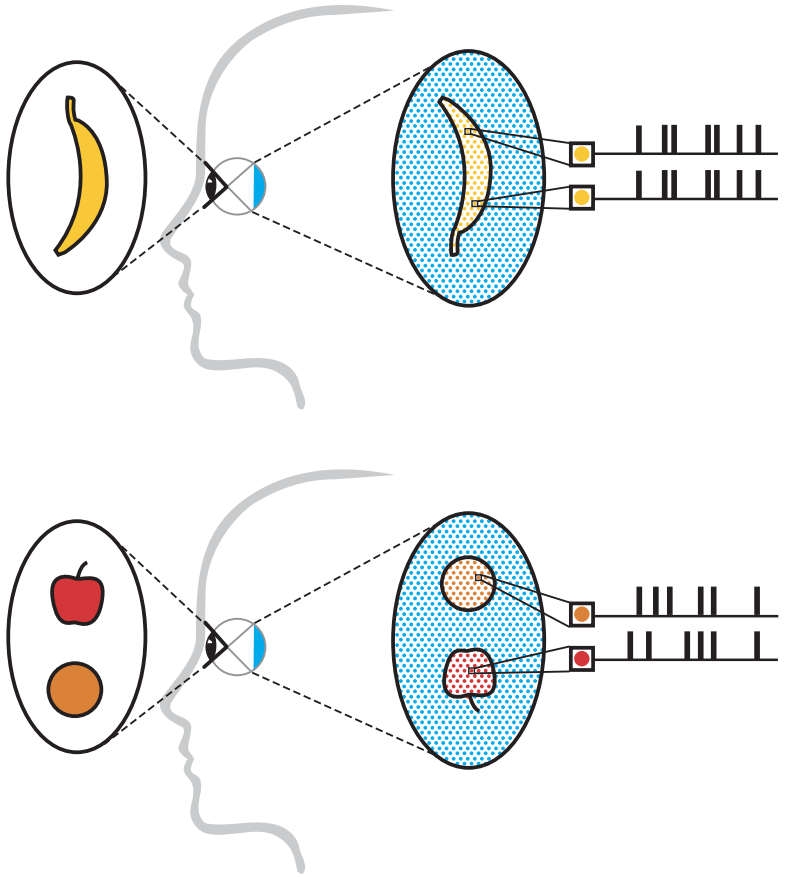

Figure 1. The retina uses synchrony to segment the visual space into objects. Top: Retinal output neurons (ganglion cells) fire synchronously when responding to the same object. Bottom: Ganglion cells fire asynchronously when responding to different objects.

Modeling the Synchronization of Ganglion Cells in the Mammalian Retina

We have constructed a detailed model of the inner mammalian retina (see Figure 2) that reproduces important aspects of the stimulus-selective synchronization between ganglion cells that has been described experimentally. Input to the model retina was conveyed by ON bipolar cells, which were driven by low-pass filtered currents (t = 2 msec) representing synaptic input from cone photoreceptors. The model bipolar cells produced excitatory postsynaptic potentials in both ganglion cells and amacrine cells (inhibitory interneurons) according to a random process. The model ganglion cells were twice as large as the bipolar cells and received inhibitory input from three different amacrine cell types encompassing three different spatial scales:

(1) small amacrine cells that were the same size as the bipolar cells;
(2) large amacrine cells, whose dendritic fields were the same size as the ganglion cells; and
(3) axon-bearing amacrine cells, whose central dendritic fields were the same size as the bipolar cells but whose long axons made synapses out to a distance of nine ganglion cell diameters, excluding a small central region corresponding to the dendritic field.

Of the three amacrine cell types in the model, only the axon-bearing amacrine cells fired spikes. All three amacrine cell types made feedforward synapses onto ganglion cells, feedback synapses onto bipolar cells, as well as synapses among themselves

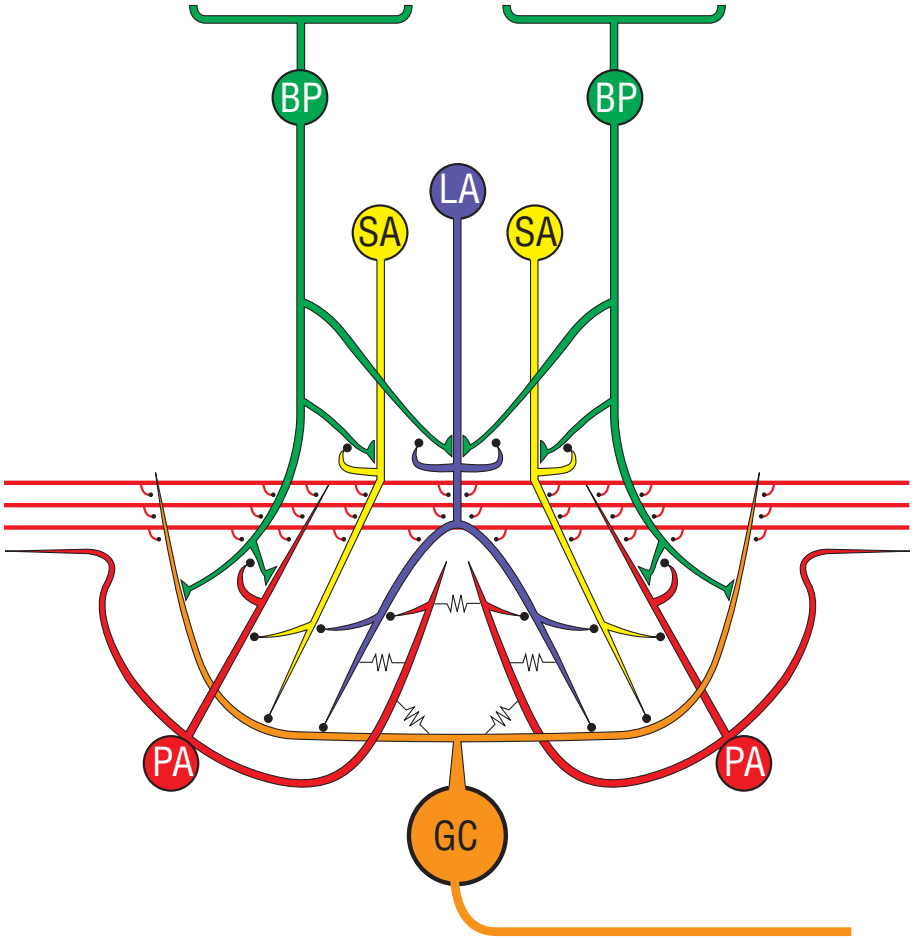

Figure 2. The model retina consisted of a 32 × 32 array of identical local processing units implemented as a torus. Input to each unit was conveyed by a 2 × 2 array of bipolar cells (BPs). The output of each unit was conveyed by the axon of a single ganglion cell (GC). Each local processing unit contained three different inhibitory interneurons, corresponding to small (SA), large (LA), and polyaxonal (PA) amacrine cells (ACs). Filled black circles are inhibitory synapses, triangular contacts excitatory synapses, and resistors gap junctions.

Responses to Small Light Spots

The data in Figure 3 plots the responses of a representative model ganglion cell to simulated light spots of increasing intensity. The illustrations on the left side of Figure 3A show the dimensions of the spot stimulus relative to the size and location of the model ganglion cell's receptive field center. The top trace on the right side of Figure 3A shows the firing activity elicited by a relatively dim spot, while the bottom trace on the right side of Figure 3A shows the firing activity produced by a relatively bright spot. While the bright spot elicits more spikes immediately following the stimulus onset, the precise temporal sequence of spikes appears more-or-less random. Electrophysiological data is therefore typically presented as in Figure 3B, which shows the peri-stimulus time histograms (PSTHs) accumulated over 20 stimulus trials for five different spot intensities. By averaging over multiple stimulus presentations, it is possible to reliably measure the time-dependent changes in firing rate produced by external stimulation.

The intensity series in Figure 3B is very similar to that recorded from ganglion cells in response to analogous stimuli,[5] thus indicating that our model network reproduces important aspects of retinal physiology. As summarized in Figure 3C, the peak firing rate following stimulus onset is strongly correlated with the intensity of the spot stimulus, while the plateau firing rate is only weakly dependent on the spot intensity. This characteristic property of ganglion cells has been termed contrast gain control[6], and in our retinal model, results from nonlinear negative feedback (see Figure 2).

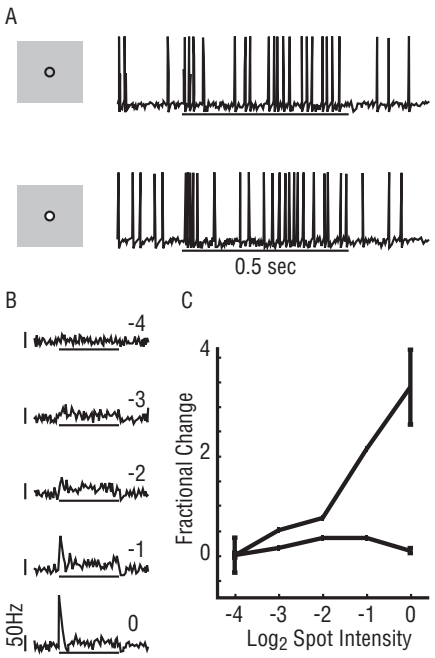

Figure 3. Responses of model ganglion cells to small spots. (A) Ganglion cell output consists of discrete, uniform pulses. A relatively dim spot (top) produces a small increase in the plateau firing rate, while a relatively bright spot (bottom) produces a large transient peak in the firing rate. (B) Peri-stimulus-time-histograms (PSTHs) plotted as a function of increasing spot intensity in $\log_2$ increments. (C) Plot of peak (solid line) and plateau (long-short dashed line) firing rate as a function of spot intensity.

Stimulus-Selective Synchronization between Model Ganglion Cells

Because the exact temporal sequence of spikes is not reproducible from trial to trial, does it then follow that the times of occurrences of individual spikes are of no significance? For many years, this was the attitude of the majority of neuroscientists. In the last decade, however, a rapidly growing body of experimental evidence suggests that the generation of discrete pulse trains provides biological neurons with a natural solution to the problem of dynamic binding and image segmentation.[7]

To investigate whether synchronous firing between the model ganglion cells was stimulus-selective, we examined the cross-correlation histograms (CCHs) between pairs of ganglion cells stimulated by two identical bars that were turned on simultaneously (Figure 4). To ensure that our analysis considered only correlations resulting from synaptic interactions, correlations resulting from stimulus coordination (dashed red traces) were subtracted. The correlations resulting from stimulus coordination were estimated by a shift-predictor, obtained by shifting one of the spike trains making up the CCH by one or more stimulus trials. The shift-predictor thus preserves correlations resulting from a common modulation of the firing rate produced by the stimulus, but eliminates correlations resulting from synaptic interactions that are not time-locked to the stimulus onset. All CCHs were normalized as a fraction of the expected baseline synchrony, given by the amplitude of the CCH at zero delay in the absence of stimulation. To further reduce the effects of stimulus coordination, only activity during the plateau portion of the response was considered. CCHs were plotted for ganglion cell pairs at opposite ends of the same bar (Figure 4b_1, upper bar; Figure 4b_3, lower bar), or at the nearest opposing tips of the two separate bars (Figure 4b_2). Significant correlations were evident between ganglion cells responding to the same bar but not between ganglion cells responding to different bars (solid black traces). Even though the ganglion cells in each pair were separated by the same distance and were stimulated identically within their receptive field centers, only those pairs that

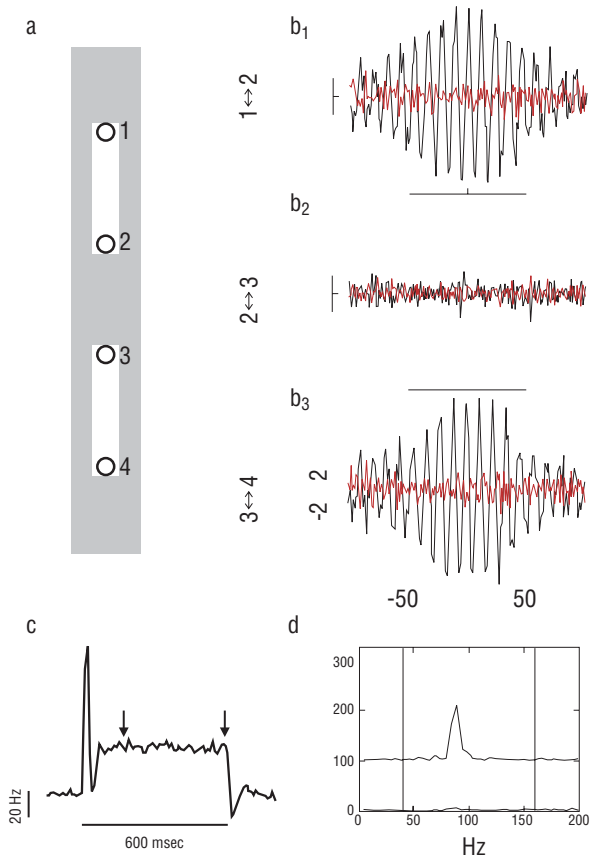

Figure 4. Stimulus-selective synchronization of ganglion cells. (a) Stimulus dimensions relative to the receptive field centers of individual ganglion cells. (b) Cross-correlation histograms (CCHs) computed during the plateau portion of the response for pairs of ganglion cells at opposite ends of the same bar or at opposing tips of separate bars. All ganglion cell pairs were separated by 7 diameters (bin size, 1 msec; scale: 100 msec, 0.5). (b_1) pair from upper bar; (b_2) pair from separate bars; (b_3) pair from lower bar. Correlations were only significant for pairs from the same bar. (c) Combined PSTH of cells 1–4. Arrows indicate plateau portion of response. (d) Power spectra of CCHs between cells stimulated by separate bars (lower trace) or by the same bar (upper trace—offset by 100).

responded to the same bar were strongly correlated. These results suggest that synchronization is able to provide a flexible label that can be used at the earliest stages of visual processing to dynamically bind together groups of elements that are syntactically related. To our knowledge, this is the first model of the vertebrate retina to account for stimulus-selective synchronization.

The CCHs between ganglion-cell pairs stimulated by the same bar show that joint firing probability oscillates as a function of the delay between the two spikes at a frequency of approximately 100 Hz, which falls within the frequency range measured experimentally.[8] There was no evidence of oscillatory activity in the combined PSTH, however (Figure 4c). This is because the 10-msec bin width used to construct the combined PSTH was too large to resolve the high-frequency oscillations evident in the CCHs. Furthermore, synaptically driven correlations only remained phase-locked to the stimulus onset for a short time (~50 msecs) and at longer times tended to average out over multiple stimulus trials. A peak in the power spectra, given by the Fourier transform of the associated CCH, is only present in the correlations between cells stimulated by the same bar, suggesting that oscillatory activity itself, as well as synchrony, can be used for segmentation (Figure 4d).

The oscillations evident in the CCHs were driven by feedback inhibition from the axon-bearing amacrine cells, as can be understood by examining the synaptic circuitry of the retinal model (Figure 2). In response to a large stimulus, ganglion cells activated neighboring axon-bearing amacrine cells by way of electrical synapses and were then hyperpolarized by the ensuing wave of axon-mediated inhibition. If the stimulus was maintained, the ganglion cells recovered simultaneously, thus setting up the next cycle of the oscillation. Axon-mediated inhibition of the neurons with receptive fields in the gap between the two stimuli furthermore helped to maintain the selectivity of the stimulus-evoked oscillations, since spikes are rapidly attenuated when passively conducted through a chain of electrical synapses.[9]

Modulation of Firing Correlations by Velocity

In addition to being highly sensitive to global properties, such as the "connectedness" between visual regions, the degree of synchrony between the model ganglion cells was also very sensitive to local stimulus properties, such as brightness, size, and velocity. An example of this is illustrated in Figure 5, which examines the correlations produced by bars moving at several different velocities. The correlations produced by slow-moving bars were similar in amplitude and frequency to those measured during the plateau portion of the step response (see Figure 4). For bars swept across the receptive field at high velocity, the change in firing rate became sharply peaked, causing the expected correlations from stimulus coordination to become non-negligible. Even at high velocities, however, the amplitude of the shift predictor was much smaller than the additional correlations resulting from synaptic interactions. The amplitude of these additional correlations increased as a function of velocity, even after correlations resulting from stimulus coordination were subtracted (Figure 5c). Measured as a fractional change from baseline, the strongest modulation was seen in the average power of the correlation function between 40–160 Hz. In contrast, the average firing rate during the same period of activity (indicated by the short tick marks under the PSTHs), was approximately constant as a function of velocity. These results demonstrate how, under some circumstances, firing correlations between many neurons can respond to changes in stimulus parameters over a greater dynamic range than the firing rates of individual cells.

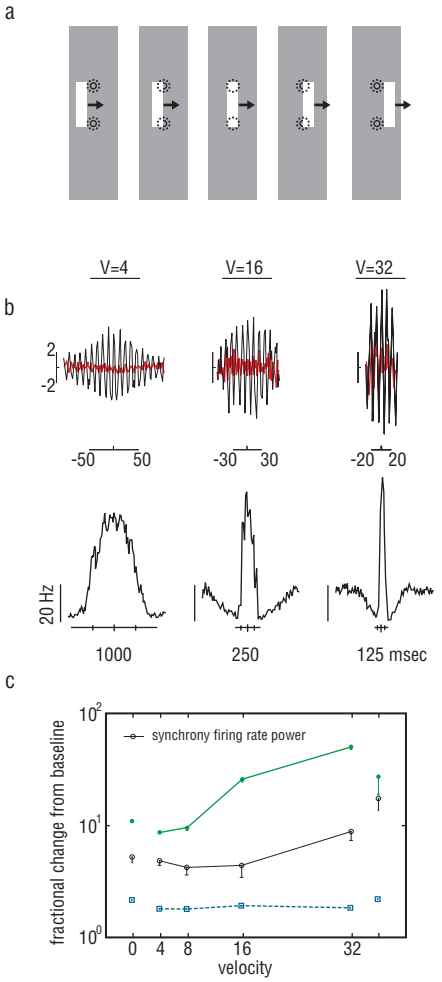

Figure 5. Firing correlations are modulated by velocity. (a) Illustration of stimulus protocol. Second and fourth panels bracket the period during which correlations were measured. (b) Ganglion cell (GC) responses at different velocities, in GC-GC distance/second, indicated at the top of each column. The top row of traces in 5b are CCHs. At high velocity, the correlations (solid black lines) become larger, while the synchrony due to stimulus coordination (solid red lines) remains relatively small. the bottom rowof 5b are combined PSTHs averaged over both recorded cells. Small vertical ticks mark measurement period. (c) The average power between 40–160 Hz (solid green line) is more strongly modulated by stimulus velocity than either the synchrony (solid black line) or the average firing rate (dashed blue line). All quantities measured as a fractional change from baseline.

Firing Correlations are Modulated by Stimulus Intensity

Additional studies were conducted to examine how firing correlations could be modulated by the intensity of a stimulus (Figure 6). A line of model ganglion cells was stimulated by a narrow bar, which was presented at a range of intensities. The combined PSTHs, averaged over all stimulated cells, showed that the average firing rate, during both the peak and plateau portions of the response, increased steadily with stimulus intensity. Firing correlations were assessed by combining the individual CCHs of all stimulated cell pairs into a massed CCH measure. The massed CCH also increased in amplitude as the stimulus intensity was raised. This was particularly true of the power spectra of the massed CCH, which developed progressively larger peaks at higher stimulus amplitudes. The location of the peak power shifted slightly towards higher frequencies as the stimulus intensity was increased, but this did not appear to be a major effect.

It is possible to compare the relative sensitivities of different measures of activity by examining their range of modulation as a function of stimulus intensity. Three measures of activity obtained during the plateau portion of the response were considered; the firing rate, the synchrony, and the average power between 40–160 Hz. Each of these measures was expressed as a fractional change from its respective baseline. With this normalization, the average power was approximately twice as sensitive as the synchrony, which in turn was approximately a factor of 2 more sensitive than the firing rate. These results suggest that significant information regarding the intensity of a stimulus can be encoded in the degree of firing correlations between the activated cells.

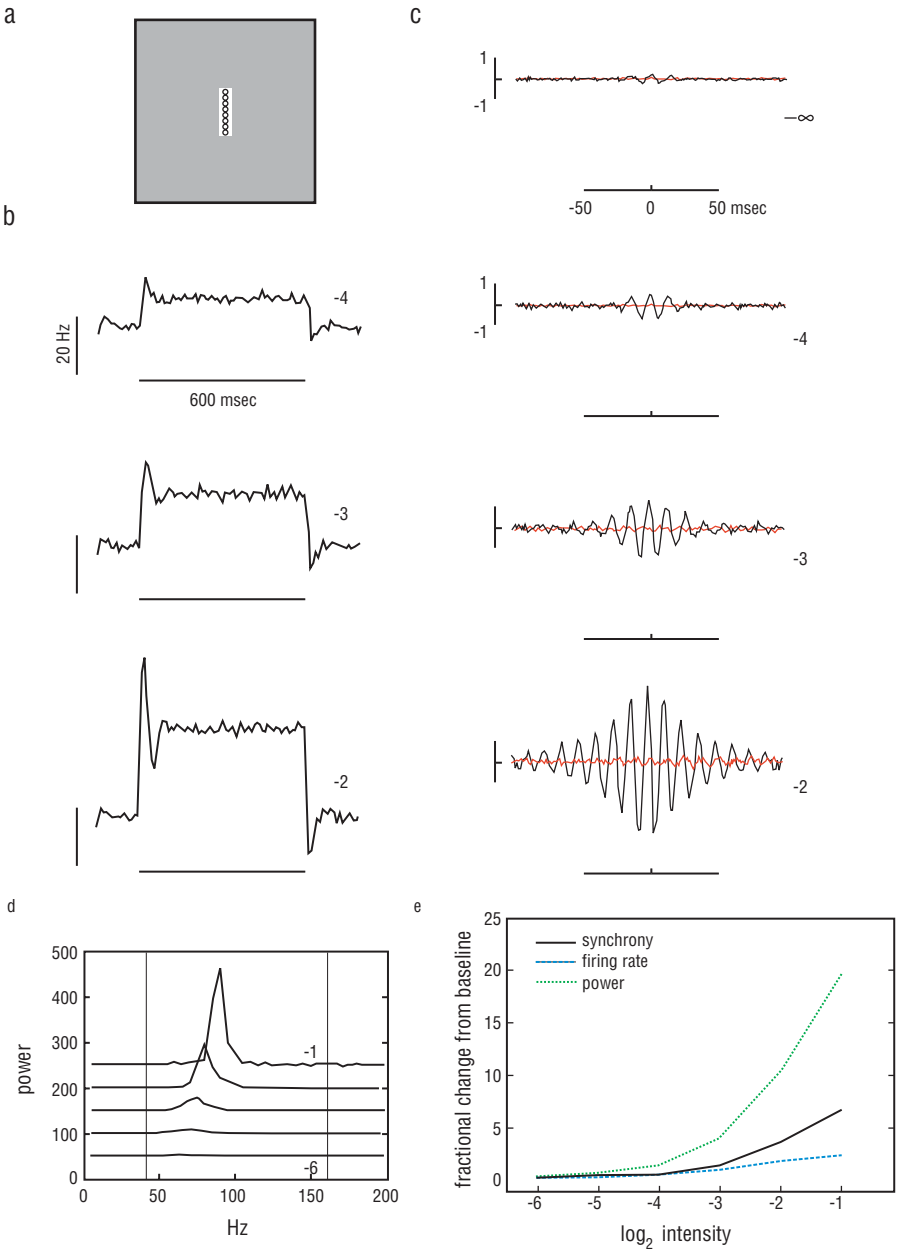

Figure 6. Firing correlations are modulated by stimulus intensity. (a) A line of eight ganglion cells was stimulated by a narrow stationary bar. (b) Combined PSTHs, averaged over all eight stimulated cells. As the $\log_2$ stimulus intensity (indicated to right of each plot) increased, firing rates increased correspondingly. (c) Massed CCHs, averaged over all stimulated cell pairs. Firing correlations also increased with stimulus intensity. (d) Power spectra of firing correlations. Increasing stimulus intensity elicited progressively more prominent peaks. Spectra at successive intensities offset vertically for clarity. (e) Sensitivity to stimulus intensity of different measures of activity. The average power in the frequency band 40–160 Hz is modulated over a greater dynamic range than either synchrony or plateau firing rate.

Using Firing Correlations to More Accurately Detect Changes in Stimulus Intensity

Our model suggests that firing correlations between retinal neurons can be modulated over a greater dynamic range than the firing rates of individual cells in response to changes in several important stimulus parameters. This finding led us to ask whether neural circuits could use the information encoded by firing correlations to more accurately resolve such stimulus parameters. To investigate this question, the spiking output of a line of ganglion cells activated by a narrow bar stimulus was fed into an ideal threshold detector (Figure 7). Each stimulus intensity was presented 200 times, and we counted the total number of threshold crossings for a 400-msec epoch during the plateau portion of the response. The average firing rate of the detector was measured at several different stimulus intensities, and the difference in detector output between all possible intensity increments is plotted in Figure 7a. For each point in the plot, the abscissa gives the final stimulus intensity, while the x-intercept of the line passing through that point yields the corresponding baseline stimulus intensity from which the difference in detector output was calculated. Each difference in detector output was normalized by its standard deviation so that the abscissa value of each point measures the reliability with which the corresponding increment in stimulus intensity could be detected. To assess the extent to which the detector utilized firing correlations between its inputs, the eight stimulated ganglion cells were replaced by independent Poisson generators that, on average, produced the same number of spikes per unit time. Overall, the normalized difference in detector output between all stimulus pairs was 45% greater using correlated input from the retinal model as compared with the Poisson control. The threshold detector illustrates how firing correlations can reliably encode information. The threshold was set to 2.5, meaning that three or more spikes were necessary to produce a detector output on any given time step. For both correlated and Poisson input, the baseline detector firing rate was very low, around 1 Hz. As shown in Figure 7b, increasing the stimulus intensity produced a significantly greater increase in the detector firing rate when the inputs were correlated as compared with the case in which the inputs were independent. This extra sensitivity was due to the fact that the threshold process is well suited for detecting rare synchronous events,[10, 11] and cortical neurons exhibit supralinear responses to synchronous inputs.[12] By encoding stimulus parameters, such as intensity, as increases in the number of such events, our results suggest that the retina may be able to transmit information more reliably along the optic nerve.

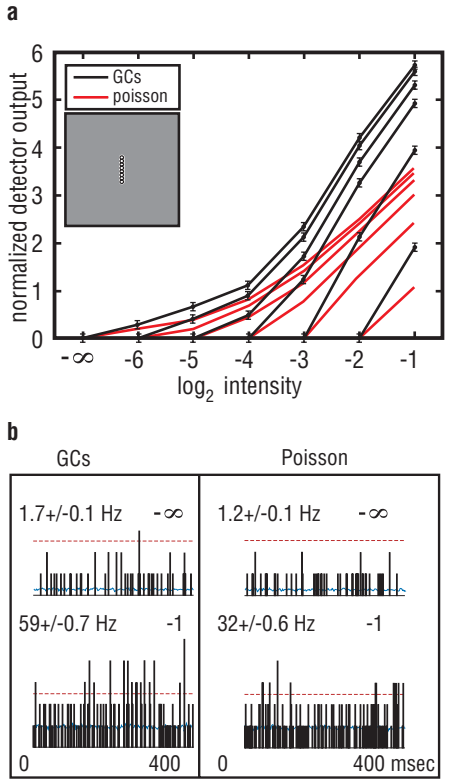

Figure 7. A threshold detector can utilize correlations between ganglion cells to better discriminate differences in stimulus intensity. (a) Incremental change in detector output in response to summed inputs from either model ganglion cells (solid black lines) or from independent Poisson generators (solid red lines). Stimuli were narrow bars of variable intensity (see inset). Each point represents the difference between the firing rate of the detector at the intensity indicated by the abscissa and the detector firing rate at the associated baseline intensity, given by the intersection of the containing line with the x-axis. Each difference was normalized by its standard deviation, thus measuring how well different intensity increments could be distinguished. Normalized detector output was 45% greater, on average, when driven by correlated input. (b) Example of the threshold detection process. Ganglion cell input is shown on the left and equivalent Poisson input on the right. The baseline activity of the detector in the absence of stimulation (top row) is very low. A stimulus with $\log_2$ intensity = -1 (bottom row) produced strong correlations between ganglion cells, resulting in a relatively large number of "rare" synchronous inputs. Dashed red line: detector threshold. Solid blue line: average firing rate per 2-msec bin.

Summary and Outlook

Our retinal model has provided insight into how patterns of retinal connectivity give rise to the high-frequency oscillations that are known to underlie the stimulus-selective synchronization of ganglion cells. We have furthermore used the model to conduct a number of computer experiments that suggest that synchrony between retinal ganglion cells not only reflects global topological properties, such as contiguity, but also encodes other important stimulus parameters, such as the brightness, size, and velocity of individual objects. In some cases, such stimulus properties were encoded more robustly by the degree of firing synchrony than by the firing rates of individual cells. Our results therefore support the contention that the rate-code hypothesis must be modified to incorporate the fundamental role that synchrony plays in how information is represented and processed in the mammalian visual system.

References

[1] P. Reinagel and R. C. Reid, "Temporal Coding of Visual Information in the Thalamus," *J Neurosci* 20(14), 5392–400 (2000).

[2] A. K. Engel, P. Konig, and W. Singer, "Direct Physiological Evidence for Scene Segmentation by Temporal Coding," *Proc Natl Acad Sci U S A* 88(20), 9136–40 (1991).

[3] M. S. Livingstone, "Oscillatory Firing and Interneuronal Correlations in Squirrel Monkey Striate Cortex," *J Neurophysiol* 75(6), 2467–85 (1996).

[4] S. Neuenschwander and W. Singer, "Long-Range Synchronization of Oscillatory Light Responses in the Cat Retina and Lateral Geniculate Nucleus," *Nature* 379(6567), 728–32 (1996).

[5] O. D. Creutzfeldt, et al., "Sensitivity Distribution and Spatial Summation within Receptive-Field Center of Retinal On-Center Ganglion Cells and Transfer Function of the Retina," *J Neurophysiol* 33(5), 654–71 (1970).

[6] R. M. Shapley and J. D. Victor, "The Effect of Contrast on the Transfer Properties of Cat Retinal Ganglion Cells," *Journal of Physiology* 285, 275–98 (1978).

[7] W. Singer and C. M. Gray, "Visual feature integration and the temporal correlation hypothesis," *Annual Review of Neuroscience* 18, 555–86 (1995).

[8] S. Neuenschwander, M. Castelo-Branco, and W. Singer, "Synchronous Oscillations in the Cat Retina," *Vision Res* 39(15), 2485–97 (1999).

[9] G. T. Kenyon and D. W. Marshak, "Gap Junctions with Amacrine Cells Provide a Feedback Pathway for Ganglion Cells within the Retina," *Proc R Soc Lond B Biol Sci* 265(1399), 919–25 (1998).

[10] G. T. Kenyon, E. E. Fetz, and R. D. Puff, "Effects of Firing Synchrony on Signal Propagation in Layered Networks," in *Advances in Neural Information Processing Systems 2*, D. S. Touretzky, Ed., (Morgan Kaufmann, San Mateo CA, 1990) pp. 141–148.

[11] G. T. Kenyon, R. D. Puff, and E. E. Fetz, "A General Diffusion Model for Analyzing the Efficacy of Synaptic Input to Threshold Neurons," *Biol Cybern* 67(2), 133–41 (1992).

[12] J. M. Alonso, W. M. Usrey, and R. C. Reid, "Precisely Correlated Firing in Cells of the Lateral Geniculate Nucleus," *Nature* 383(6603), 815–9 (1996).